Low-complexity serially-concatenated coding for the deep space optical channel

Bruce Moision, Jon Hamkins[1]
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109-8099
e-mail: {bruce.moision,jon.hamkins}@jpl.nasa.gov

I. Introduction

NASA is developing optical links to support deep space communication at data rates on the order of 100 Mbit/second. Power efficient signaling is achieved by modulating the data using M-ary pulse-position-modulation (PPM). For certain lasers and detectors, the optimal PPM order is high–$M \geq 256$.

We illustrate performance 0.5–1.0 dB from capacity via the iterative decoding of a serial concatenation of a short constraint length convolutional code and coded PPM through a bit interleaver. To realize the gains of the iterative decoding requires likelihoods to be computed and stored for each PPM symbol. High data rates, large values of M and large interleavers can make the likelihood computation and storage prohibitively expensive.

To reduce the complexity of iterative decoding, we propose to compute and store only a subset of the channel likelihoods. We show that this can be done while suffering a negligible loss in performance. The complexity of implementing the forward-backward algorithm is also reduced when partial likelihoods are retained.

II. Partial Statistics

Suppose only P of each M observations, as well as their indices $\boldsymbol{\mathcal{I}}$, are made available to the receiver. We model the constraint as shown in Figure 1, where $\mathbf{c}, \mathbf{n}, \mathbf{y}$ are codeword, noise and received M-vectors, respectively, ϕ denotes the rule for choosing the P samples, $\phi : \mathbf{y} \rightarrow (\mathbf{y}_{\boldsymbol{\mathcal{I}}}, \boldsymbol{\mathcal{I}})$, $\boldsymbol{\mathcal{I}} \subset \{1, \ldots, M\}, |\boldsymbol{\mathcal{I}}| = P$ and $\mathbf{y}_{\boldsymbol{\mathcal{I}}}$ is the vector of retained samples $\mathbf{y}_i, i \in \boldsymbol{\mathcal{I}}$. The mapping ϕ reduces the dimensionality of the received information, which we expect to severely degrade performance for small P. However, by allowing $\boldsymbol{\mathcal{I}}$ to be a function of the observation $\mathbf{y}$, we will see that the reduction may be implemented with negligible loss.

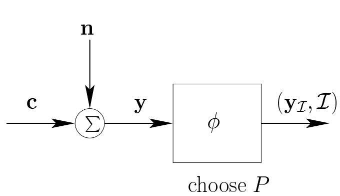

Figure 1: Constrained storage channel model

[1]This work was supported by the IPN Technology Program and the Caltech Presidents Fund and performed at the Jet Propulsion Laboratory, California Institute of Technology under contract with the National Aeronautics and Space Administration.

The conditional likelihoods

$$p_{\phi(\mathbf{y})|\mathbf{c}}(\phi(y)|c) = p_{\mathbf{y}_{\mathcal{I}}|\mathbf{c}}(y_{\mathcal{I}}|c) p_{\boldsymbol{\mathcal{I}}|\mathbf{c},\mathbf{y}_{\mathcal{I}}}(\mathcal{I}|c, y_{\mathcal{I}})$$

where we explicitly distinguish random variables $\mathbf{y}, \mathbf{c}, \boldsymbol{\mathcal{I}}$ from their realizations $y, c, \mathcal{I}$, are a sufficient statistic for estimation of $\mathbf{c}$ given $\phi(\mathbf{y})$. The term $p_{\mathbf{y}_{\mathcal{I}}|\mathbf{c}}$ is the likelihood that would result if the input were mapped to a P dimensional constellation and the term $p_{\boldsymbol{\mathcal{I}}|\mathbf{c},\mathbf{y}_{\mathcal{I}}}$ is an adjustment to reflect the outcome of the decision.

Figure 2 illustrates performance for a case with $M = 256$, a 4096-bit interleaver, and 8 iterations, on an AWGN channel. We see 0.1 dB degradation when 1/64 of the likelihoods are kept, and negligible degradation when 1/32 are kept.

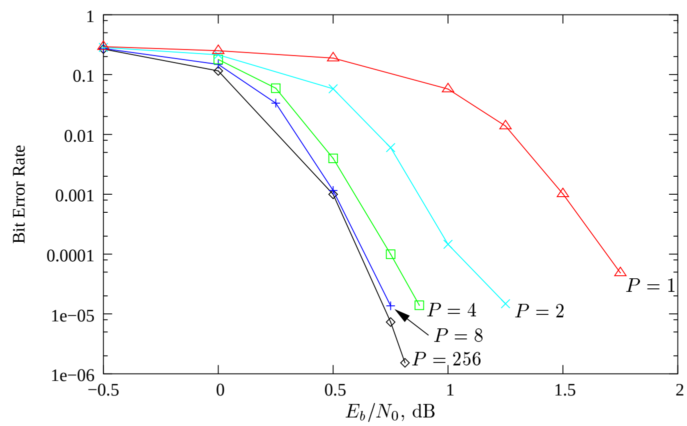

Figure 2: Performance with partial statistics, $M = 256$.

In addition to storing the floating point values, we must save the addresses of the P largest values. This yields

$$\frac{\text{storage for } P \text{ likelihoods}}{\text{storage for } M \text{ likelihoods}} = \frac{P}{M}\left(1 + \frac{\log_2 M}{f}\right),$$

where f is the number of bits used to represent fixed or floating-point values.

III. Partial trellis

With partial statistics, there are only $P + 1$ distinct channel likelihoods $p_{\phi(\mathbf{y})|\mathbf{c}}(\phi(y)|c(e))$. One can take advantage of this and use a reduced complexity time-varying trellis. Let $\mathcal{J}_k$ be the collection of parallel edges in the full trellis at time k with the same channel likelihood. Form a *partial* trellis by replacing the edges in $\mathcal{J}_k$ with a single edge with the same initial state and terminal state. Maximum aposteriori decoding may be executed on the partial trellis. The partial trellis will have $|\mathcal{V}|$ states and no more than $|\mathcal{V}|(P + |\mathcal{V}|)$ edges, where $|\mathcal{V}|$ is the number of states on the full trellis.